CHAPTER

2

PLANNING YOUR SYSTEM

After reading this chapter and completing the exercises, you will be able to:

- Discuss hardware issues related to installing Linux
- Collect hardware and related information from several sources
- Prepare space for Linux on a system running Windows

In the previous chapter you learned about the development of Linux, beginning with Linus Torvalds' school project and continuing with contributions from developers around the world. You learned about the license under which Linux is distributed and how that license affects the way people build businesses and careers on Linux.

In this chapter you learn about how Linux uses the hardware resources in your computer system. This information will help you plan your installation of Linux, so that you can make the right choices for your system.

PREPARING TO INSTALL LINUX

Most of the computer systems that you have used or purchased yourself probably had an operating system installed. As a result, you didn't have to think about how the operating system was installed on the computer's hard disk or how the operating system was configured to use other parts of the computer, such as the mouse, the video display capabilities, and so forth.

Although you can purchase computers with Linux preinstalled from companies like VA Linux Systems (*www.valinux.com*), you will learn a great deal about your computer system and about how Linux operates by installing the operating system yourself. Before you begin the installation of Linux in Chapter 3, you must decide a few things about how you want to install Linux. You must also compile information about your computer hardware so that you can answer questions that arise during the installation process.

As a general rule, Linux is not more difficult to install than other operating systems such as Windows. The difference is that you rarely install Windows—it comes preinstalled on most

systems that you buy. But if you ever have to install Windows from scratch, you may find that it can be more difficult to install than Linux.

Understanding Computer Hardware

The design of the Linux operating system is based on the concept of the computer as a collection of devices. Information is stored on a hard disk device, output is written to a video card device, input is read from a keyboard device, and so forth. Linux must be configured to use all of the devices on the computer system in order to function correctly. Unfortunately, this is not always an easy task, as you will learn in the later sections of this chapter.

Linux can be used on many types of computers, including those that use different types of microprocessors (CPUs), such as Alpha, SPARC, and PowerPC. The discussion in this chapter is devoted solely to computers using Intel and Intel-compatible microprocessors—in other words, standard PCs.

Different types of devices on a computer communicate with the operating system in different ways. This means that you must gather an array of information about your computer system before you can install Linux on that system. The section "Creating a System Inventory," later in this chapter, explains how to locate the information you need.

As you work with computer hardware, you should be familiar with a few common terms that describe that hardware. Space or capacity on a computer system is measured in bytes. A **byte** is enough space to store one character. Because computers store many characters, space is commonly measured in megabytes, abbreviated MB. One **megabyte** is 1,048,576 bytes, or enough space to store roughly 1 million characters. Another common term is gigabyte, abbreviated as GB. One **gigabyte** is 1024 MB, or roughly enough space to store 1 billion characters.

Storing Information

The electronic memory of a computer is called **random access memory**, or RAM. Information in RAM is only available when the computer is turned on. When you turn the computer off, everything stored in RAM is lost. RAM is normally measured in MB, with most computers having from 16 MB to 128 MB of RAM.

Another electronic component in a computer, called **read-only memory**, or ROM, stores information about how the computer starts and how the computer's devices are configured. Like RAM, ROM is stored on a computer chip, but information in ROM is not lost when the computer is turned off; it is permanent, or nonvolatile. One of the key things stored in ROM is the Basic Input/Output System, or **BIOS**, which provides instructions to the operating system for using the devices on the computer. The BIOS itself cannot be changed—it is permanent. But the ROM that contains the BIOS also contains parameters that control parts of the computer configuration. You can change these settings using a special BIOS utility.

Information in the BIOS might include the setting for the computer's clock, information about each disk drive, the settings used by the modem, a start-up password, and many other details.

You can usually access a menu to view and reconfigure settings related to the BIOS by pressing a key while the computer is starting. To find out which key, watch the screen during start-up for a message similar to "Press Del for Setup." You then have a few seconds during the system start-up to press the specified key, at which point the computer will display the BIOS configuration menus rather than starting the operating system. You will see example BIOS configuration menus later in this chapter.

The **hard disk** is a magnetic storage space for data, such as the operating system and data files that you create. You can think of magnetic storage as being like the stripe on the back of a credit card, except that a hard disk in a typical new computer holds from 3 billion to 10 billion characters (3 GB to 10 GB of data). Hard disk storage is not permanent—you can make changes to information on the hard disk—but it is also not volatile, meaning that storage on the hard disk remains intact when the computer is turned off. When you turn on your computer, the BIOS loads information from the hard disk into RAM for regular operations. When the computer is switched off, the information in RAM is discarded, but the data on the hard disk remains, ready to be reloaded the next time the computer is turned on. Table 2-1 highlights the differences between RAM, ROM (where the BIOS is stored), and the hard disk.

Table 2-1 Computer Storage Components

Component	Permanent (cannot be changed by a computer user)	Volatile (disappears when the power is turned off)
RAM	No	Yes
ROM	Yes	No
Hard disk space	No	No

A computer can have multiple hard disks. Each one is considered a separate device. In an operating system like Windows, two disk drives would be called C: and D:. The naming scheme is different in Linux, as you will see in Chapter 3.

When you see a reference to Windows in this chapter, the reference applies to Windows 95, Windows 98, Windows 2000, and Windows NT, unless otherwise noted.

Device Communications

Many computer devices communicate with the microprocessor and software programs via interrupt requests. An **interrupt request**, or IRQ, is a numbered channel of communication through which a device can ask the operating system to perform a specific action. A PC

has a limited number of IRQs, ranging from 0 to 15. To configure some devices, such as Ethernet networking cards, you may have to know the IRQ that the device uses by default, or you may have to change the IRQ that it uses. Later in this chapter you will learn how to determine the IRQ used by a device.

Once a device sends an IRQ signal to the microprocessor, the device and the CPU can communicate data and status information between them. This is done using two methods: direct memory access and memory-mapped input/output.

A **direct memory access (DMA) channel** allows a device to read and write directly to the computer's RAM, without going through the microprocessor first. DMA allows a device such as a sound card to read and write information to memory much more quickly than if the microprocessor were involved in each data transaction.

Not all devices use DMA, however. Most use **memory-mapped input/output** (memory-mapped I/O)—a technique that assigns a range of memory addresses in the computer as a place for a device to send and receive data. Essentially, this kind of memory works like a post office box. The device places data in a specific memory location; then software programs retrieve the data from that memory location and place new data there. The device then retrieves the data placed by the software. For a device to use memory-mapped I/O, it must have a memory address assigned to it. Most devices can use several different addresses, so you can configure your computer to avoid a conflict between two devices that try to use the same address. The memory addresses used by a device for memory-mapped I/O are referred to as the **I/O ports** for the device. Some devices, such as sound cards, require both a DMA channel and an I/O port, but this is unusual.

To refer to an I/O port you use a computer memory address. These addresses use a different numbering system because of the way bytes store information. You may have heard of binary numbers, in which everything is represented by zeros and ones. Another commonly used system is called hexadecimal. **Hexadecimal** numbering is a base-16 counting system. It uses the letters *A* through *F* (usually capitalized) to count the numbers 10 through 15. Using hexadecimal (often called hex) numbers is strange at first. The important point is that when you encounter strange numbers that contain letters, be assured that they are simply hexadecimal numbers. Be certain to write them down and enter them in Linux using the format that you see them in, including all the letters.

Hexadecimal numbers are often written with a prefix of 0x to identify them as base-16 numbers. For example, you might see the number 0x220 used as an I/O port address. The 0x indicates that this is a hexadecimal number. You don't need to convert it to normal (decimal, or base-10) numbers. Just use it in Linux as it's written.

The range of IRQ numbers from 0 to 15 can all be represented by a single hexadecimal digit. The possible IRQ numbers in hex are 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, F.

Hard Disk Devices

The hard disk in a computer must communicate with the microprocessor using an electronic interface that controls how data is sent and received. Two interfaces are used for hard disks in PCs: IDE (integrated drive electronics) and SCSI (small computer systems interface).

IDE is a low-cost, easy-to-manage interface used by most new computers to connect hard disks and CD-ROMs to the CPU. Standard new PCs always include one IDE hard disk. An **IDE controller** card is a device that handles communication between the hard disk and the microprocessor. (Sometimes the IDE controller is integrated into the system board rather than being a separate expansion card.) To use an IDE controller card to connect to the CPU, the hard disk must be compatible with the IDE interface. You will hear such hard disks referred to as IDE hard disks. Each IDE controller can be connected to two hard disks. Many computers come with two IDE controllers. The second IDE controller is used to communicate with a CD-ROM drive. Figure 2-1 shows how two IDE controllers, each with a separate cable, can be connected to hard disk and CD-ROM devices.

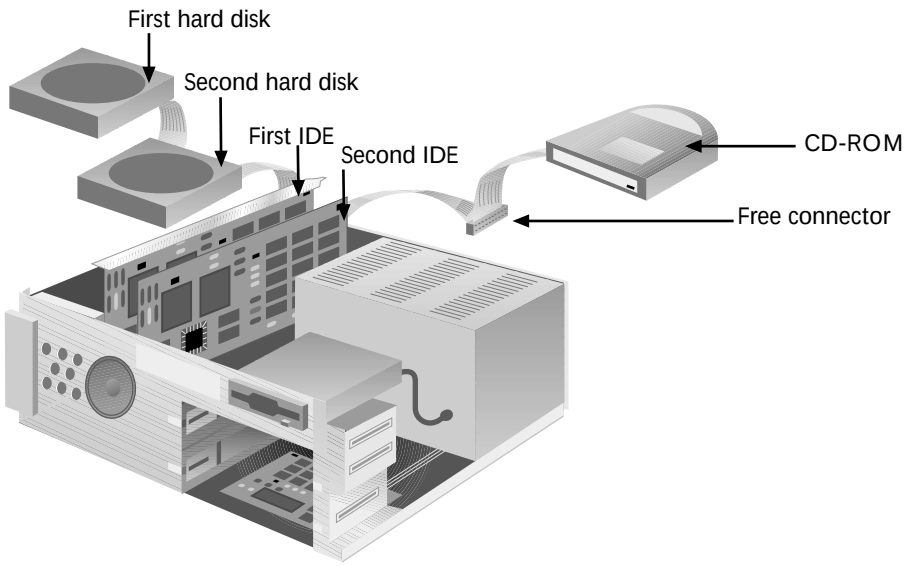

Figure 2-1 Hard disk and CD-ROM devices on multiple IDE interfaces

SCSI is a high-performance interface used by many types of devices to connect to a computer. SCSI is much faster than IDE, meaning that the flow of data between a hard disk or other device and the CPU is much more rapid. But SCSI devices are also much more expensive than IDE devices. As with IDE, a SCSI controller card provides the connection between

SCSI-compatible devices and the CPU of the computer. A single SCSI controller card can connect up to eight devices, each linked by a cable, as illustrated in Figure 2-2.

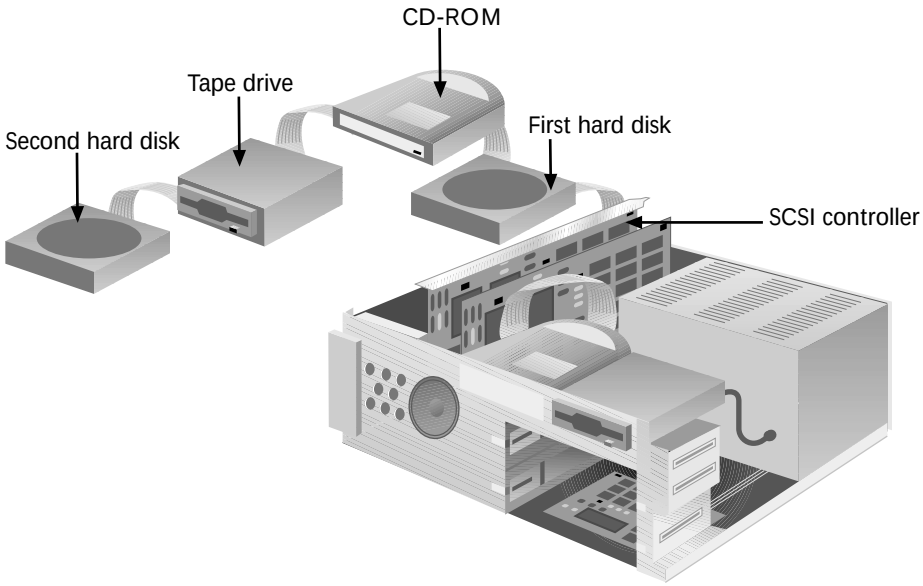

Figure 2-2 Multiple SCSI devices are linked together by cables

Supported Linux Hardware

Linux works with thousands of different hardware devices, from very old proprietary CD-ROM drives to the latest high-speed networking cards. But not all devices work with Linux. For example, Linux cannot accommodate the WinModem devices included with some new PCs.

You are likely to have trouble using Linux with very recently released hardware (such as a brand-new video card with many new features) and with very old equipment that uses standards that have since been discarded. The more information you have about the device in question, the more likely you can configure Linux to use the device, even if complete support for all of the device's features is not available.

Before installing Linux, you may want to check the Web site of your Linux vendor. The Web site may include a list of all supported hardware for that distribution. The hardware supported by all distributions is basically identical, so you can review the hardware compatibility list from any site as a starting place (for example, visit *http://www.calderasystems.com/support/hardware/*). Different versions of the Linux kernel, however, support slightly different hardware.

Understanding Networking

Many Linux systems are connected to networks so that they can communicate and share resources with other users connected to the network. In order to configure networking on

Linux, you will need some configuration information, as described later in this chapter. A basic knowledge of networking concepts will help you understand how to use that configuration information.

To communicate with each other effectively, computer systems use a **protocol**, which is an organized pattern of signals or words. Linux networking involves many different protocols. One protocol provides basic communication between network adapters. Another lets Web browsers communicate with Web servers. Altogether, Linux supports dozens of networking protocols, each one designed for a different purpose. Some of the most important protocols are discussed in the following sections.

IP Networking

The most important protocol in relation to Linux networking is called the Internet Protocol, or IP. **IP** is a networking protocol used to send packets of information across a network connection. IP is the basis on which most Linux networking is built. In order to use IP networking, each computer on the network must be assigned an identifying number, called an **IP address**. Each packet of data sent across the network includes two IP addresses: an address to identify the computer that the packet is coming from and one to identify the computer that the packet is going to.

IP addresses consist of four numbers, each separated by a period. For example, 207.29.12.1 is an IP address; so is 192.168.100.15. Each of the four numbers can be a value from 0 to 255. IP addresses are used all over the world, so to avoid the confusion that arises from duplicate IP addresses, you must use an IP address that has been assigned to the computer on which you install Linux. You may have to do some checking to find out what authority assigns the IP address for your computer. For instance, your Internet Service Provider (ISP) may have assigned a range of numbers to your system administrator, who would in turn assign a number to your computer as needed.

All of the IP addresses on a single network are related; for example, their first three numbers may be the same. Several special IP addresses are associated with setting up Linux networking so that packets can be passed around the network. These special IP addresses are listed here:

- A **network mask** tells the networking system in Linux how to identify IP numbers that are part of the local network, as opposed to IP numbers that are assigned to computers outside the local network.
- A **broadcast address** identifies a special IP address that will send a packet of data to all computers on the local network.
- A **gateway address** identifies the computer that can send packets of data outside the local network, to the Internet or to other networks in an organization.

A good Linux installation program can calculate some of the different IP addresses just described using the IP address of your computer. Still, your system administrator will probably give you each of the IP addresses needed for your computer and network.

Some networks use a special protocol designed to avoid the need to set up and track IP addresses on a network. This protocol, called the **Dynamic Host Configuration Protocol**

(DHCP), allows a computer to obtain an IP address dynamically from a network server at the time the computer is turned on. This can make it much easier for the system administrator to manage networking on a large system. A range of IP addresses and all the special addresses for the network are stored on a central server. Each computer is then configured to look for the DHCP server when it boots.

Rather than specifying an IP address, you may choose to use DHCP as you install Linux, depending on the configuration of your local network. You cannot use DHCP unless a DHCP server has been configured on your network. In addition, not all types of network services are well served by a computer that uses DHCP. For example, if you want to create a Web server on your Linux system, you should use an assigned IP address rather than relying on DHCP.

Domain Names and Hostnames

Transmitting data across a network would be difficult if you had to remember the IP address for every computer you wanted to access. To simplify matters, you can make use of a system of human-readable names for computers and networks. Each of these names is mapped to (is associated with) a specific IP address. As described in the next section, a special network server does the work of translating computer and network names into IP addresses, and vice versa.

A name assigned to a network is called a **domain name**. (Domain names are actually more all-encompassing than this definition implies, but equating a domain to a network is generally correct.) Some examples of domain names are *compaq.com*, *linux.org*, and *nasa.gov*. Domain names within a large organization may be longer than this. For example, within Compaq, you may find domains called *marketing.compaq.com*, *sales.compaq.com*, and *research.compaq.com*. The domain name is not limited to two words, but the ending portion of the domain name will always be one of the standard top-level (most generalized) domain names. Table 2-2 shows a few of the top-level domain names you are likely to see. Not all of the top-level domain names are listed, because each nation has a separate top-level domain name.

Table 2-2 Top-Level Domain Names

Name	Description
.com	Commercial/business entities
.org	Noncommercial organizations
.net	Organizations whose work relates to the Internet
.edu	Educational institutions, usually colleges and universities in the United States
.gov	U.S. government organizations
.mil	U.S. military organizations
.us	Generalized top-level domain for networks in the United States; used mostly for local governments and schools
.de	Top-level domain of Germany (Deutschland)
.uk	Top-level domain of the United Kingdom

Each computer in a domain is assigned a name. Computers involved in networking are often called **hosts**; so the name of a computer on the network is referred to as the computer's hostname. The **hostname** is a single word used to name a computer. Your system administrator may let you choose a hostname, or may assign one, such as lab13, or training01.

The hostname is combined with the domain name to create a **fully qualified domain name (FQDN)**. An example of an FQDN would be *lab13.myschool.edu*. Using the longer domain names given previously as examples, a sample FQDN might be *taco.research.compaq.com*. Web page addresses take the form of FQDNs. For example, the Web address *www.ibm.com* consists of a hostname (*www*) and a domain name (*ibm.com*). How can you tell that this is not simply a long domain name? By using a Web browser to successfully access a single system using that FQDN. (You can't tell by looking at the name.) In theory, IBM could have a domain named *www.ibm.com*. A host in that domain with a hostname such as *florida* would yield a FQDN of *florida.www.ibm.com*.

Domain Name Service (DNS)

The **Domain Name Service (DNS)** is the special network service devoted to the task of mapping human-readable domain names and hostnames to the IP addresses of specific networks and computers. A **DNS server** is the computer that actually performs this conversion. The process works like this: When you enter the address *www.ibm.com* in a Web browser, the browser must send a network packet to a DNS server asking for the IP address of *www.ibm.com*. Once that address is returned by the DNS server, the Web browser can establish a connection to the IBM Web server using the IP address (204.146.80.99 in this case).

Generally, you must provide the IP address of a DNS server as you configure Linux networking so that domain names and FQDNs can be converted to IP addresses. If you work on a small network that is not connected to the Internet, you can use a list of hostnames and IP addresses on your Linux system rather than a DNS server. Most networks employ a DNS server, however.

Creating a Shared System

Installing Linux on a system that already contains another operating system, such as Windows 98, allows you to experiment with Linux and take advantage of its features while relying on another operating system to support other needs, such as running applications that are not available for Linux. A shared system, also called a **dual-boot system**, is one that allows you to choose which operating system to start each time you boot (turn on) your computer. You can actually have numerous operating systems installed on one computer, not just two. The name dual-boot is commonly used to refer to any system with more than one operating system installed.

A program called a **boot manager** lets you select an operating system each time you boot the computer. **LILO (Linux Loader)** is the boot manager included with Linux. You can also use other commercial boot manager programs such as System Commander or BootMagic. Installing the LILO boot manager is part of every Linux installation, as you will learn in Chapter 3.

To create a dual-boot system, you must decide where on the computer's hard disks each operating system will reside. Two basic options are available:

- Store each operating system on a separate hard disk.
- Store multiple operating systems on the same hard disk.

The first option, using a different hard disk for each operating system, is preferable because all data from the different operating systems remains separate and installation is simpler. Of course, not many computer systems come with multiple hard disks, so this chapter focuses on how to store multiple operating systems on the same hard disk.

To use multiple operating systems on a single hard disk, you must decide how much space you will need for each operating system. The first operating system installed on the computer (often a version of Windows) probably takes up a lot of space already. You can determine the amount of hard disk space used on a Windows system as follows:

1. Double-click the **My Computer** icon on the Windows Desktop.
2. Right-click the icon for the hard disk. (This will normally have a name, followed by (C:), to indicate that it is drive C:.) A shortcut menu opens.
3. Click **Properties**. The Properties dialog box opens.
4. View the information on the General tab of the Properties dialog box, which is shown in Figure 2-3. This window shows the amount of used space and free space, as well as other information that will be covered later in this chapter.

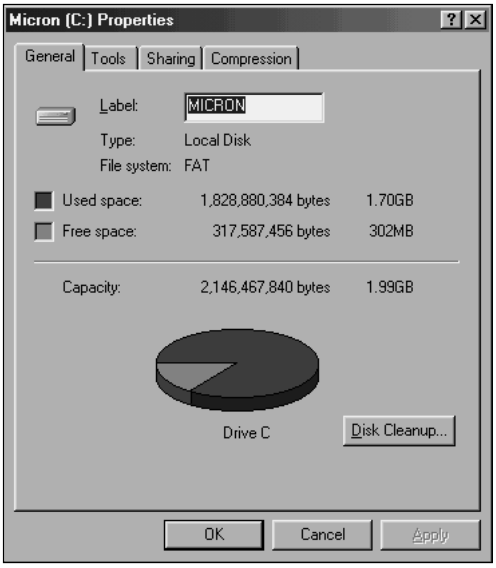

Figure 2-3 Hard disk information in the Windows Properties dialog box

5. Click **Cancel** to close the Properties dialog box.

Once you know how much space is available on the hard disk where Windows is stored, you can consider how much of the available free space you could use for your Linux operating system.

As with other operating systems, the exact components you need to install depend on what you intend to do with Linux. The more components you add, the more space you need on the hard disk. Each Linux distribution offers different standard installation options. Some distributions (such as Red Hat) let you select a Desktop system or a Server system; others (such as Caldera OpenLinux) let you choose a Web Server, Desktop System, Business Server, or Development Workstation (for use in developing software). Many distributions (including SuSE, Red Hat, and others) let you choose the exact software packages that you want to install. This is more time-consuming than just selecting a designation such as Web Server and letting the installation program choose what to install based on that selection. The sizes of a few standard options that you might see as you install different Linux distributions are shown in Table 2-3.

Table 2-3 Typical Linux Installation Options

Type of installation	Typical hard disk space required	Comments
Minimal installation without graphics support	40 MB	A specialized Linux system can use as little as 1 MB, but such specialized systems are not part of most Linux distributions.
Standard installation	400 MB to 600 MB	This generally includes a graphical system and numerous utilities.
Web server	300 MB	This includes all the networking utilities to run a Web server, but may not include graphical support.
Development workstation	500 MB	This allows you to develop Linux software.
Full (everything included with the distribution)	1.5 GB to 4 GB, depending on how many CD-ROMs are included with the distribution	

If you don't want to make a lot of decisions, choose one of the installation types shown in Table 2-3. At the other end of the spectrum, you could specify each and every option, a time-consuming process. As a compromise between these two options, you might consider a custom installation that lets you choose categories of packages. For example, in Red Hat Linux, you can choose a Custom installation option, then select packages from a list of about 20 categories.

Use the guidelines in Table 2-3 along with any documentation you have for your Linux distribution to determine how much space you need in order to install Linux.

In addition to space for the operating system, you must consider how much hard disk space you will want to have available *after* installing Linux to hold documents and other files that you create as you work in Linux. Depending on what you will do with Linux, you might need a lot of space for database files, Web server documents, programs that you download, or graphics files that you create. Add this amount to the amount of disk space for the operating system itself.

Don't let all the installation options presented so far confuse you. If you are just starting out with Linux, try to have about 1 GB of hard disk space to install Linux and work with Linux programs. This should be sufficient to get you started.

Graphical Systems

Setting up the video hardware that provides a graphical interface for Linux is undoubtedly the most challenging part of Linux configuration. A few old standards for video hardware exist, and these are still useful in configuring Linux. Newer video hardware uses proprietary technology that requires special software for each new product that becomes available. What's more, while video card vendors can be counted on to develop device drivers for Microsoft Windows, few vendors provide software support for Linux. At the same time, the vendor may refuse to release technical information about the video card, making it very difficult for Linux developers to provide their own support for the latest video cards.

On a positive note, when technical information is available, the graphical system in Linux can make full use of the features of all video hardware, without being restricted to the small number of standard modes that the vendor's Windows software supports.

The graphical environment in Linux is provided by the X Window System. Software from the **XFree86 Project** (another free software project) adds the X Window System to Linux. You can also use commercial software to add the X Window System to Linux, but the XFree86 software is included with all distributions of Linux and will be the focus of the graphical installation described in Chapter 3.

You can read about the two major commercial X products for Linux by visiting Xi Graphics at *www.xig.com* and MetroLink at *www.metrolink.com*.

The Web page for the XFree86 Project contains a complete list of video cards supported by their software. New video cards are added regularly as new versions of XFree86 are released. Video card manufacturers are slowly beginning to recognize the fact that having support in XFree86 software helps to sell more video cards, and so they are becoming more open about providing technical information to the free software community.

In addition to knowing which video card your computer system uses, you might also be able to locate the name of the video chipset used. This is helpful because sometimes when the video card model is not listed in the Linux configuration options, the chipset used on that card is listed. The chipset information can then be used to configure the graphical environment.

Monitors and the Graphical System

Because the X Window System is flexible enough to use all the features of a video card, it may be able to produce better graphics than your monitor can display. This can cause a newer, powerful video card to send signals that damage older monitors. For example, if you have a video card that supports a high-resolution display of 1240 × 1024 pixels, but your monitor is not capable of displaying that resolution, the signal from the video card might damage the monitor.

The solution to this potential problem is to configure Linux with the correct information about your monitor's capabilities. Thus, Linux will calculate what the monitor can support and block any signals that the monitor cannot handle. In the example just given, a suitable lower resolution mode would be selected to display graphics.

Three specifications identify a monitor's capabilities:

- The number of screen redraws per second (called the refresh rate). This number is expressed in hertz, or Hz. New monitors often have a range of refresh rates, such as 50 to 90.
- The number of lines per second that can be redrawn. This number is expressed in kilohertz, or KHz.
- The clock speed of the monitor's electronics, which determines indirectly how many colors the monitor can support. This number is expressed in megahertz, or MHz.

When installing most Linux distributions, you will be asked for the first of these numbers, the Hz value.

Creating a System Inventory

Before you start installing Linux, you should gather all the information about your hardware that you might need to complete the installation. Table 2-4 lists the information you may need, with example values for each and a blank column where you can fill in your values. The following sections explain how to locate the necessary information.

Although the information in Table 2-4 is useful, it may not be strictly necessary in order to install Linux. The latest Linux installation programs attempt to make installation as easy as possible, automatically detecting many hardware components. This is especially true for newer computers with standard equipment such as IDE hard disks. Nevertheless, the information in Table 2-4 can be valuable *after* you have installed Linux and need to update configurations for the LILO boot manager, the graphical system, or other parts of the Linux operating system.

Table 2-4 Useful Hardware Information for Installing Linux

Hardware information	Sample value	Your system
Amount of RAM	64 MB	
Hard disk interface type	IDE	
Hard disk size	4.3 GB	
Serial port used by modem	COM1 (first serial port)	
IRQ used by modem	4	
Printer port	LPT1 (first parallel port)	
CD-ROM interface type	IDE	
Type of mouse	Microsoft serial	
Port used by mouse	COM2 (second serial port)	
Monitor make and model	NEC Multisync C400	
Monitor scan rate range	50–90	
Video card make and model	Starfighter AGP	
Video card chipset	Intel i740	
Amount of video RAM	8 MB	
SCSI card make and model (if you have one in your system)	Adaptec AIC-7850	
Network adapter make and model	3Com Fast Etherlink XL 10/100	
The IRQ used by the network adapter	11	
Sound card make and model	SoundBlaster	
The IRQ, DMA channel, and memory addresses used by the sound card	10, 3, 220	

In addition to information about your hardware, you may need to find the networking information listed in Table 2-5.

Table 2-5 Information Needed to Set Up Linux Networking

Network information	Sample value	Your system
Domain name	xmission.com	
Hostname	brighton	
DHCP server address	207.49.12.1	
Static IP address	192.168.100.13	
Network mask	255.255.255.0	
Broadcast address	192.168.100.255	
Gateway address	192.168.100.1	
DNS server address	192.168.100.1	
Secondary DNS server address	207.29.12.2	

Installations that use a network connection to access the installation data may also require that you provide the information in Table 2-6.

Table 2-6 Information Needed for a Networked Installation

Network installation information	Sample value	Your system
NFS or SMB server IP address	192.168.100.4	
Path to the installation files on the remote server	`/mnt/cdrom`	

You can get the information described in these tables from a variety of sources. It's probably easiest to ask a technician or system administrator who is familiar with the hardware you are working on. However, if you employ the methods described in the following sections, you will learn more about how your system functions and be better prepared to become a system administrator yourself.

Finding the Manuals

Nearly every computer system includes some type of printed manual that describes how to set up and use the system. Unfortunately, this documentation is usually filed in some forgotten corner minutes after the system is up and running. Although manuals seem to be getting smaller every year, yours probably contains at least a few key specifications for your system. Locate your computer's manual and do the following:

- Look in the index under the name of the hardware component you want to learn more about.
- Check the Table of Contents for a Specifications or a Troubleshooting section.
- Review the first few pages of the section about setting up the system, which may contain details about the hardware.

Most computers come with separate manuals for each component; each of these manuals is provided by a different manufacturer. (Sometimes these manuals are tucked inside a sleeve with a CD-ROM full of software for the device.) Look for separate manuals on the following items:

- Main system (the CPU)
- Monitor
- Mouse
- Video card
- Modem
- CD-ROM drive

If you can't locate the printed manuals, try going online. Visit the Web site for the manufacturer of your computer. Search under Products, Technical Support, or a related topic. Locating technical information on a huge Web site can be time-consuming, but this information is generally free. A sample Web page showing the specifications for a Toshiba laptop is shown in Figure 2-4.

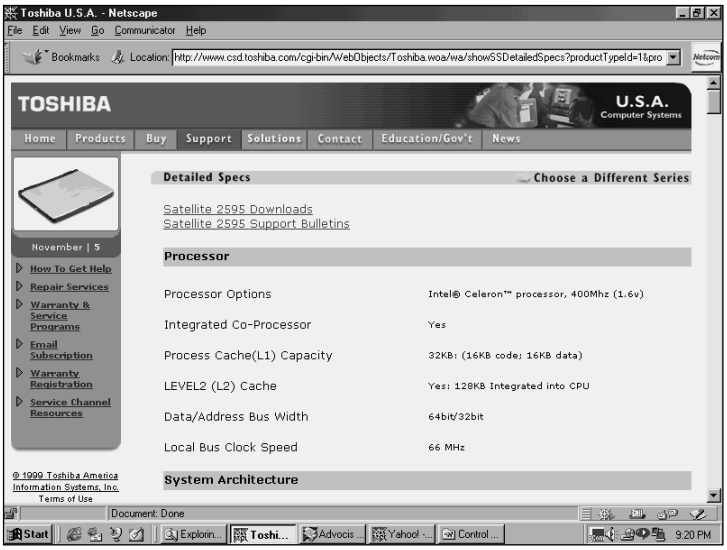

Figure 2-4 A sample Web-based specifications sheet

Reviewing BIOS Settings

Sometimes you can find information that is not evident in the printed or online documentation by reviewing the BIOS of the computer itself. Different systems use various methods of accessing the BIOS configuration menus, where you can learn about system status and devices. An on-screen message normally explains how to enter the Setup or BIOS information screen when you first boot the computer. Depending on your system, you might be

asked to press F2, Del, Esc, or some other key to enter the BIOS menus. Review the system documentation if you don't see a message explaining how to enter the BIOS menus.

Exact steps for exploring the BIOS menus cannot be given because each manufacturer may use a different interface for configuring the BIOS. A sample screen for BIOS configuration is shown in Figure 2-5.

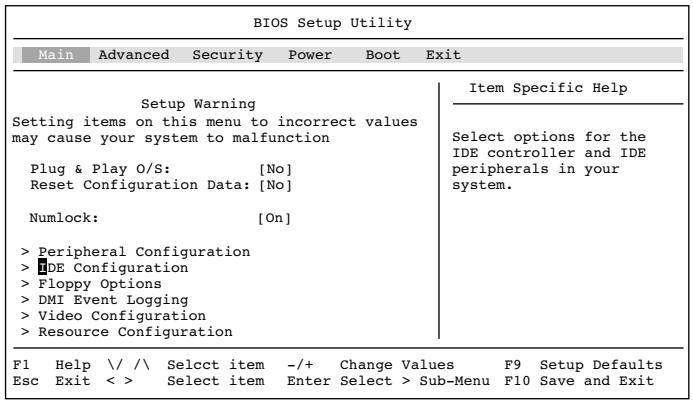

Figure 2-5 BIOS configuration menus

Keep the following points in mind as you examine the BIOS menus for hardware information:

- Some BIOS configurations include a number of options that are not relevant to installing Linux, such as how long to wait before entering a low-power mode and whether to include a power-on password. You can ignore these options and focus on locating the information in Table 2-4.
- Different levels of information may be provided in different menus. The details in Table 2-4 are often in the Advanced portion of the BIOS configuration, because most users don't need to know this information to use a system that is already up and running.
- You can make minor changes to the settings in the BIOS and save them as you exit the BIOS setup utility. Don't do this unless you have studied the system documentation and are familiar with the features you are altering. Watch carefully as you exit the BIOS menus to make sure you do not accidentally choose a Save and Exit option.

You may find any of the following pieces of information as you review your BIOS settings:

- Amount of RAM on the system.
- Hard disk interface type and size.
- Serial ports available on the system and the IRQ used by each one.

- IRQ numbers used by other devices that might conflict with devices such as a sound card or network adapter. (You usually can't see the IRQ used by these devices in the BIOS menus.)
- The SCSI card make and model.

Without accessing the BIOS, you can use a related method for locating system information: simply watch the screen carefully as the system starts. Many devices, especially SCSI cards and video cards, print identification messages to the screen as they are initialized at system boot time. You may have to power the system off and on several times to read the messages, but you can often gain much useful information from these small "advertisements" that the card prints on screen.

Studying Microsoft Windows

If your computer is already running a Windows operating system, you have an advantage in preparing for installation: the Windows operating system has already collected all of the hardware configuration information for you. Before installing Linux, you can start your computer in Windows and write down all the configuration information you'll need to use when installing Linux.

Device information in Windows is available from the Windows Control Panel, which you can view by clicking Start, pointing to Settings, and then double-clicking Control Panel. You can also double-click the My Computer icon on the Desktop, and then double-click the Control Panel icon. The Windows Control Panel (in large icon view) is shown in Figure 2-6. The icons on each Windows system vary slightly based on the hardware and software that you have installed; so your screen may vary slightly from the figure.

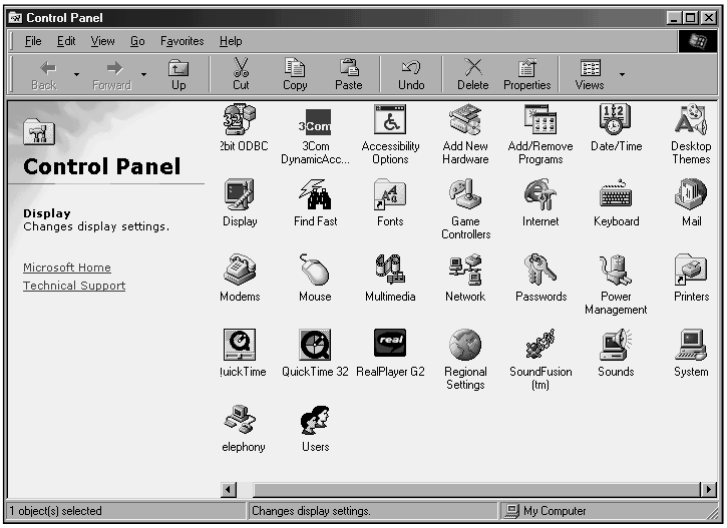

Figure 2-6 The Windows Control Panel

You can view device information in the Control Panel by double-clicking the System icon. Within the System Properties dialog box, the Device Manager tab shows a list of all the device categories on your system, such as disk drives, network adapters, and monitors. When you click the plus sign (+) to the left of an item, a more detailed list of information appears for that item. Figure 2-7 shows the Device Manager tab with several of the subsections opened for a more detailed view. As with the Control Panel, the information in this figure will differ slightly from what you see because different computer systems have different devices installed.

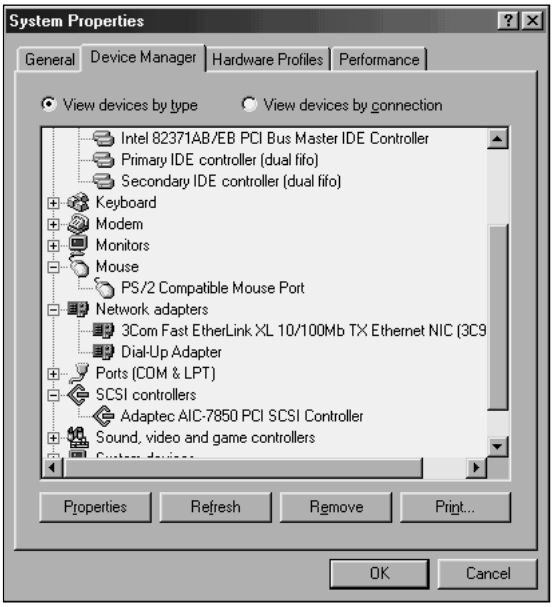

Figure 2-7 The Device Manager tab in the System Properties dialog box

In the list of devices, you may be able to see the manufacturer and model name for the device used on your system. Figure 2-7 shows a 3COM Fast Etherlink XL device as a network adapter and an Adaptec AIC-7850 as a SCSI controller. You can research additional information about your system's hardware by selecting a hardware device and choosing the Properties button at the bottom of the dialog box. (You can also double-click on any item in the list.) A separate dialog box will then appear for that device. Most of these device-specific dialog boxes include a tab labeled Resources.

The Resources tab for many devices shows the interrupt request number, the DMA channel, and the I/O port used by the device (labeled Input/Output Range in the dialog box). You can copy all of the information that you feel will be useful (based on what you have read so far) in setting up the Linux system. Not all of the information is necessary. For example, the keyboard device lists an IRQ and an input/output range, but the keyboard is automatically configured in Linux, so you won't need that information. Refer to Table 2-4 to review the

hardware information you may need. Figure 2-8 shows a sample Resources tab for a device in the Device Manager.

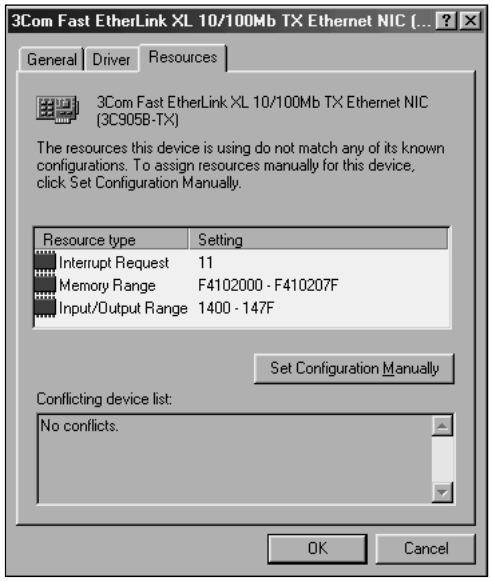

Figure 2-8 The Resources tab in the Properties dialog box of a selected device

You can also use Windows to find details about your networking configuration. Because a computer cannot run both Windows and Linux at the same time, you can normally use the same networking information for Linux that you use for Windows.

Within the Windows Control Panel, the Network icon leads to information about how networking is set up in Windows. The Identification tab in the Network dialog box includes a Computer name field with a name for your system. This name is associated with Windows networking, but you may want to use this as your hostname when you install Linux. Ask your system administrator for advice.

The Configuration tab of the Network dialog lists many different Windows networking components. One item is labeled TCP/IP, followed by the name of your networking device (this is usually an Ethernet networking adapter). Figure 2-9 shows a Network dialog box with the TCP/IP item selected.

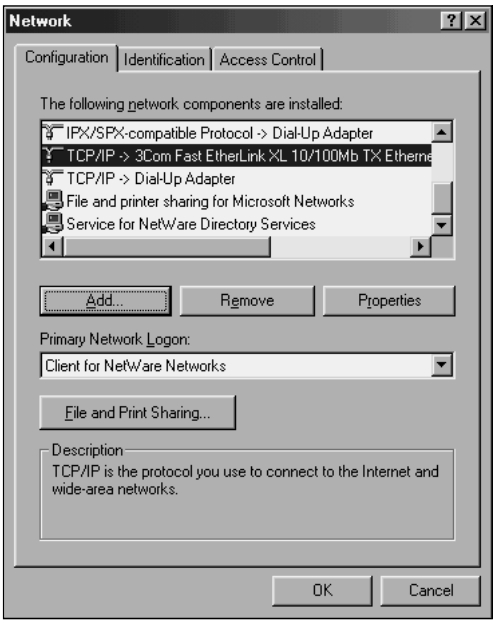

Figure 2-9 The Network dialog box in Windows

With the TCP/IP item selected, click the Properties button. This opens the TCP/IP Properties dialog box. In this dialog box you can collect the following information:

- On the IP Address tab, you see the IP Address field and the Subnet Mask field (which refers to the Network Mask field as described previously). The Obtain IP address automatically item indicates that the Windows system uses DHCP. If this option is selected in Windows, you can try to select DHCP as you install Linux.
- On the Gateway tab, the first item in the section, labeled Installed gateways provides the IP address that you should use as the gateway address when installing Linux.
- On the DNS Configuration tab, the DNS Server Search Order list contains at least one IP address. If multiple addresses are shown, you should make note of them all. The Host and Domain names above the DNS Server Search Order list are the hostname and domain name used for Windows networking. Unless directed otherwise, you should use these as the hostname and domain name when you install Linux. If the Disable DNS option is selected, the information in this dialog box may not be valid, or it may not be shown as described. In this case, ask your system administrator or ISP for the DNS-related information mentioned here.

Figure 2-10 shows the DNS Configuration tab in the TCP/IP Properties dialog box. The values shown differ on each system, based on the domain name and DNS servers used for each network.

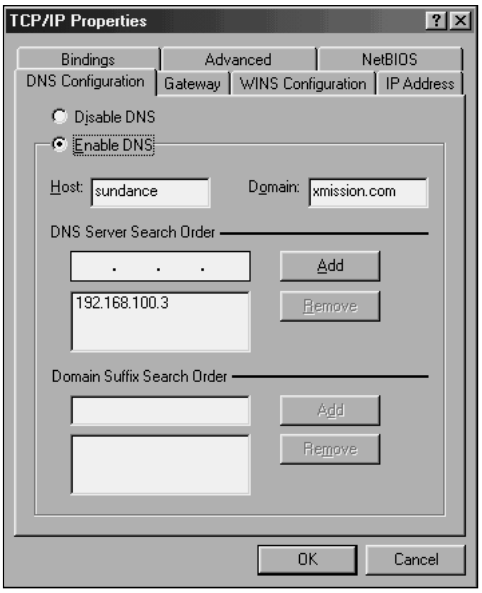

Figure 2-10 The DNS Configuration tab in the TCP/IP Properties dialog box

Exit all of the Windows dialog boxes described here by choosing the Cancel button. This prevents you from accidentally altering your Windows settings and causing problems with your hardware or networking configurations.

Asking Networking Questions

The only way to obtain some network settings is to ask the person (or organization) who assigns that setting. For instance, you cannot simply choose a network address, nor is the network address of your gateway or other servers something that you can guess. To get most of the networking information listed in Table 2-5 and Table 2-6, you will need to ask the authority who originally set up your network connection—that is, you need to consult either your system administrator or ISP.

In some cases, you will be allowed to select a hostname for your Linux installation. You can choose any brief name that you want for the hostname. System administrators often use a pattern of names for setting up multiple computers, such as the names of pets, cities, or something similar. The names are arbitrary, but you must provide the name to the person who manages the name server on your network so that e-mail and other services can be directed to your computer.

PREPARING YOUR HARD DISK

You must install Linux in a dedicated partition on the hard disk. A **partition** is a distinct area of a hard disk that has been prepared to store a particular type of data. For example, a computer that only contains Windows has one partition on its hard disk. That partition is marked as containing Windows data. To install Linux you must prepare another partition that is marked as containing Linux data. You must create a Linux partition whether you plan to establish a dual-boot system or install only Linux on the computer.

You can think of a partition as an empty space on the disk with a label attached to it. Before a partition can hold information, it must be formatted with a particular file system type. The term **file system** refers to the arrangement of information on a device such as a hard disk. The organization usually takes the form of files and directories. In theory, you could create a partition that was marked as holding a certain file system type, then format that partition with a different file system type than the partition specified. In practice, however, the operating system prevents this sort of mistake.

Figure 2-11 shows how two partitions might be arranged on a hard disk, with marks indicating the file system type and a file system format inside each partition. The default file system type for Linux is called **`ext2`** (extended file system version 2). The default file system type for Windows 98 systems is called **FAT32** (32-bit File Allocation Table). Windows NT uses a file system called **NTFS** (NT file system).

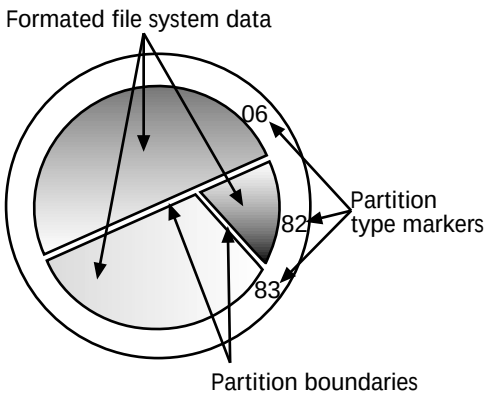

Figure 2-11 Partitions and file systems on a hard disk

Among all of the partitions on a hard disk, one can be marked as the active partition, or the bootable partition. The **active partition** is the one that the BIOS assumes contains the operating system to be started. The BIOS always attempts to start the boot manager program located at the beginning of the active partition.

In Chapter 3 you will learn in detail about the tools you can use to create Linux partitions. To use these tools, you must have free space available on your hard disk. For computers that will be dual-boot systems, you must prepare the free space for Linux before you start the Linux installation. For computer systems that will have only Linux installed, the partitions

can be created during the installation. The information in the following sections will help you understand how to use the partitioning tools discussed in Chapter 3 and how to create free space on a Windows partition.

Booting the System

When you turn on a computer, the BIOS initializes the devices on the system, then starts the boot manager located on a small area of the first hard disk called the **Master Boot Record**, or MBR. The MBR contains a small program that decides how to start an operating system. Normally, the MBR passes control of the computer to the boot record of the active partition. The **boot record** is a small area on each partition that contains a boot manager program for launching the operating system on that partition. For partitions containing Linux, the boot record contains a copy of the LILO boot manager.

The boot manager can include instructions to pass control to other partitions, such as a partition containing a Windows operating system. This creates a dual-boot system in which you choose which operating system to start each time you turn on the computer.

Figure 2-12 shows how a hard disk is arranged to include an MBR and a boot record on multiple partitions, with one partition being marked as active. Sometimes the boot record is called the boot sector, or the root sector, of the partition.

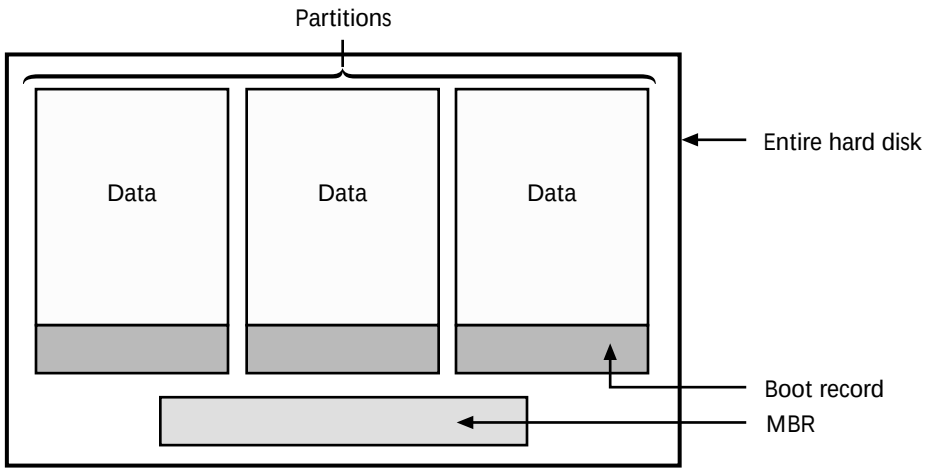

Figure 2-12 A hard disk with an MBR and boot records on each partition

Hard Disk Geometry

The operating system stored on a hard disk arranges information according to the file system used. Each operating system has a default file system. But the hard disk itself is designed to store information according to the physical characteristics of the hard disk. Sometimes you need to know something about the disk drive layout to configure the system properly.

A hard disk is composed of multiple flat platters that hold magnetic data. These platters are stacked together, with small reading devices moving between the platters to read the data on each platter. Each concentric circle on a platter is called a **track**. When you format a hard disk, each track is divided into multiple sectors. A **sector** is a unit of data storage on a hard disk. Normally a sector contains 512 bytes. Sectors are often grouped together into larger units called clusters or blocks. A default hard disk **block** in Linux is a unit of hard disk space that contains 1024 bytes, or two sectors. Figure 2-13 shows a single platter of a hard disk with the tracks, sectors, and blocks illustrated.

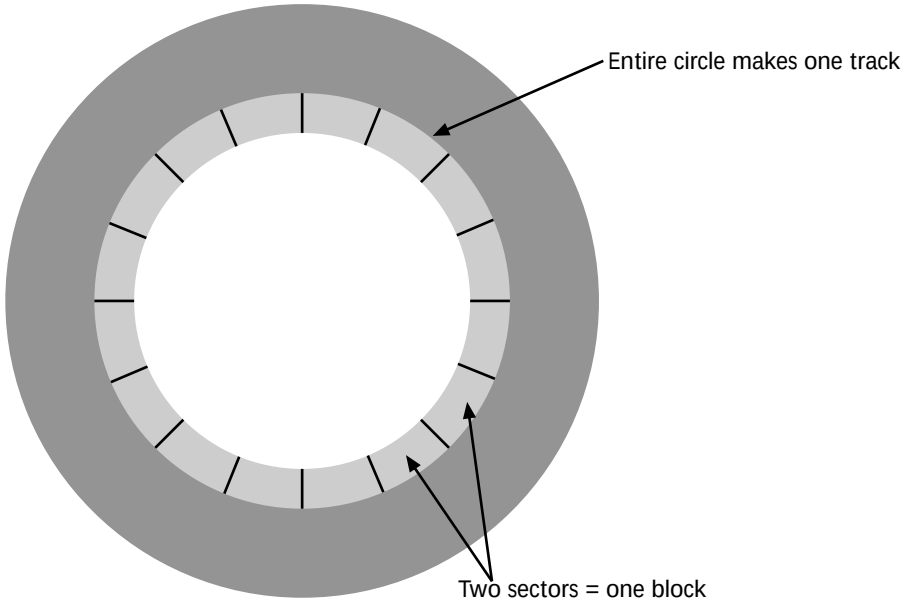

Figure 2-13 Tracks, blocks, and sectors on a hard disk platter

Taken together, all of the tracks at the same position on each platter are called a cylinder. Put another way, a **cylinder** is a set of tracks at the same location on all the platters of a hard disk. Figure 2-14 illustrates this concept.

Note Figures 2-13 and 2-14 are somewhat deceptive. Keep in mind that a hard disk often has over 1000 cylinders, or concentric tracks, on each platter. This is an important point. Hard disks with many cylinders can cause special installation problems for Linux because of limitations in the BIOS of many computers. The next section explains this problem.

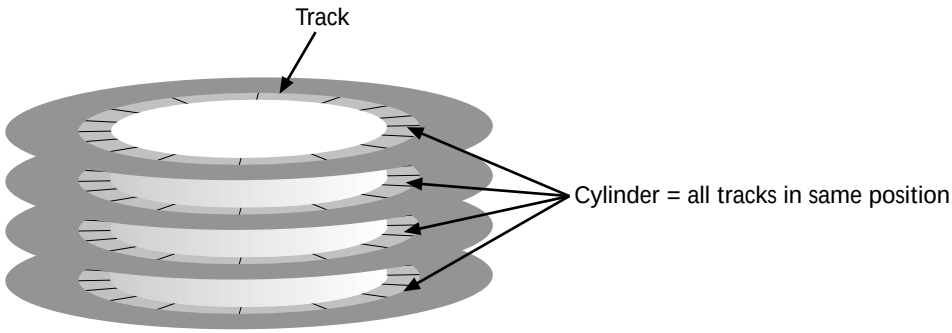

Figure 2-14 A cylinder on a hard disk

Using Big Disks

When installing Linux, the terms big disk or large disk refer to hard disks with more than 1024 cylinders. When you define partitions on your hard disk, you must define them on cylinder boundaries. That is, you cannot define a single cylinder as being partly in one partition and partly in another partition. Because of a limitation in the BIOS programming in many computers, the MBR cannot send control to a partition that includes cylinder numbers higher than 1024. The BIOS simply cannot handle larger numbers.

For example, suppose you define the partitions on your hard disk like this:

Windows: Cylinders 1–700
Linux: Cylinders 701–1400

Even if you mark the Linux partition as the active partition, the BIOS cannot pass control to the boot manager stored on the Linux partition because some of the cylinder numbers that make up the partition are too large for the BIOS to address. To deal with this strange problem, you sometimes must define your Linux system to include two partitions. First you define one small partition using smaller cylinder numbers; this partition can boot the operating system. Then you define a second, larger Linux partition that includes space on the hard disk beyond the 1024-cylinder boundary. Another operating system may reside in the space between the two Linux partitions. For example:

Linux start-up: Cylinders 1–200
Windows: Cylinders 700–1200
Linux: Cylinders 701–2000

The number of cylinders that you use for each partition depends on the amount of space you require for each operating system, but it also varies based on the total number of cylinders available on the hard disk. Most Linux installation programs will help you overcome the so-called big disk problem as you install Linux.

Swap Partitions

The Linux operating system and the files that you create in Linux are stored in a partition of type ext2. In addition to this partition, you must create a separate partition used as a swap partition for the Linux kernel's virtual memory. **Virtual memory** is a special area of the hard disk that the operating system treats as if it were RAM, storing programs there temporarily when they are not being used. The **swap partition** (also called swap space) is the designated area used as virtual memory by the Linux kernel. For example, you might start several programs in Linux, but if one of the programs is sitting idle and the RAM memory that the program uses is needed by another program, the Linux kernel will copy part of the program or its data to the swap partition. When the program is activated again, the program and data are moved back to RAM so they can continue operating.

Using a separate partition controlled by the Linux kernel allows the virtual memory feature to work very quickly, without interfering with other hard disk operations that read from or write to the `ext2` file system. A swap partition is normally from 32 MB to 256 MB in size, though it can be smaller or larger depending on how much hard disk space you have and how busy the Linux system will be. Some Linux systems limit a swap partition to 128 MB.

You create a swap partition just as you create a regular Linux partition. The partition is marked with an indicator that it will be used for Linux virtual memory.

Preparing a Shared Hard Disk

Dual-boot systems require that you either have a separate hard disk for Linux or that you use some of the hard disk space currently occupied by Windows to install Linux. If you choose the latter approach, you need to become familiar with **FIPS**—a program that lets you create two separate partitions from your existing Windows partition. During the Linux installation, you can delete the second, empty Windows partition and replace it with a Linux partition.

Before using the FIPS program, you should back up any important data on your Windows system by copying it to disks, tape, or to a network server. FIPS is considered stable and safe, but altering hard disk partition information always puts your data at risk.

The FIPS program is a DOS program that is included with every Linux distribution, usually in a separate directory called `utils`, `dostools`, or something similar. Your Linux documentation will tell you where it is located on the CD-ROM. In order to use FIPS, you must first arrange all the data on the Windows partition so that it is together, leaving a contiguous area of free space at one end of the hard disk. In other words, you must defragment the Windows system. **Defragmenting** is a procedure that arranges each file in Windows so that the parts of the file are next to each other on the hard disk (as opposed to the parts of a file

being fragmented, or spread across the hard disk). When you defragment a Windows system, all of the files are placed at the beginning of the hard disk. Figure 2-15 shows conceptually how a Windows partition is arranged before and after the defragmenting operation.

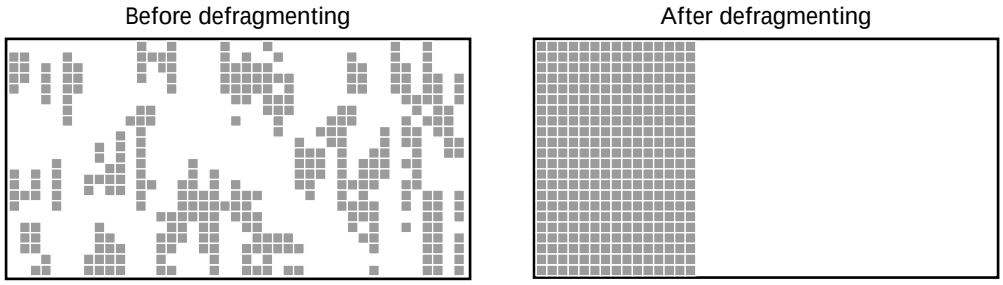

Figure 2-15 A Windows partition before and after defragmenting

Windows 98 includes a utility for defragmenting your Windows partitions. To defragment a Windows partition within Windows 98:

1. Double-click the **My Computer** icon.
2. Locate the hard disk on which you will later install Linux (this is normally C: in Windows).
3. Right-click the icon for that hard disk. A shortcut menu appears.
4. Click **Properties**. The Properties dialog box for the selected hard drive opens.
5. Click the **Tools** tab to display the options shown in Figure 2-16. The last item on the Tools tab is the Defragment utility.

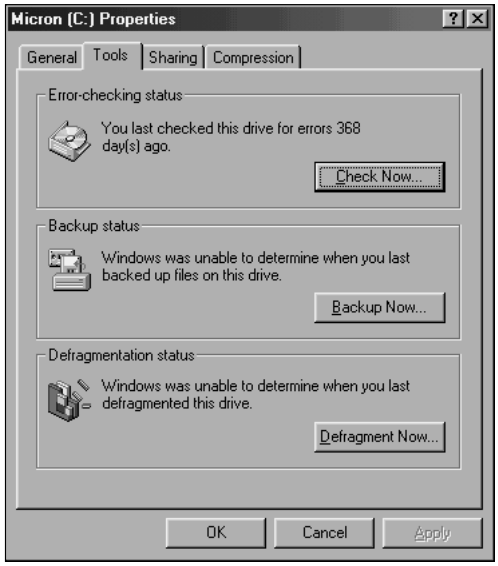

Figure 2-16 The Tools tab for a Windows hard disk

6. Click the **Defragment Now** button. The Defragmentation program runs on the selected hard disk. You see a progress indicator as the program works.
7. Click **Show Details** in the progress indicator window to see how the Defragmentation program is rearranging data on the hard disk. Figure 2-17 shows the detailed screen as the program works.

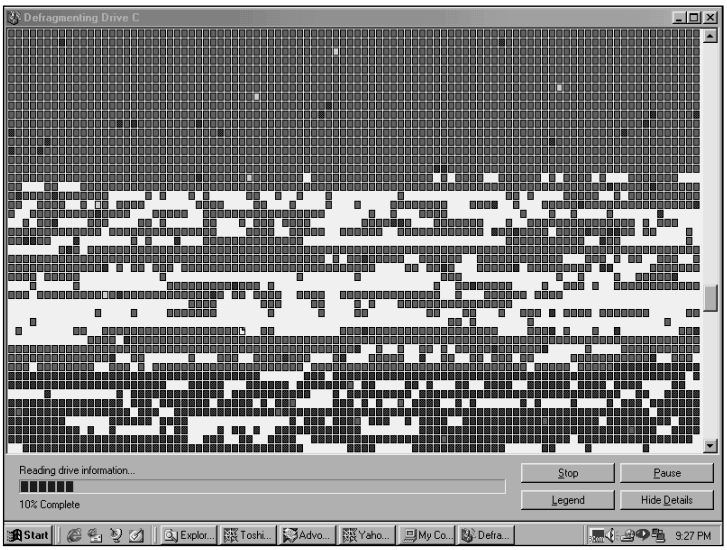

Figure 2-17 The Defragmentation details window

8. When the program finishes, click **Yes** to exit.

With the free space arranged on your Windows partition, you can run the FIPS program to split the Windows partition into two pieces. To run FIPS, follow these steps:

1. Obtain a floppy disk to use as a backup for your disk information.
2. Obtain a second floppy disk to hold the FIPS program.
3. Locate Fips.exe on the Linux installation CD (or among the Linux files that you have downloaded from the Internet). The CD may contain several copies of FIPS.exe in differnet folders. Use any copy that you find.
4. Copy Fips.exe to a floppy disk.
5. Click the **Start** button in Windows, click **Shut Down**, then click **Restart in MS-DOS mode**. After a moment, the DOS C:\ prompt appears.
6. Start the FIPS program using the command **FIPS** preceded by the drive letter for your floppy drive (for example, A:\FIPS). The CD-ROM drive is normally not available within the DOS mode, which is why you copied the FIPS program to a floppy disk.
7. Read the messages on screen, and then press a key to continue.

8. Using the floppy disk mentioned in Step 1, back up your disk information by following the instructions on screen. (Keep this disk until you have finished installing Linux and made certain that you can start both Windows and Linux.)
9. If you have multiple disk drives, select the one you want FIPS to alter.
10. The partition table for the hard disk is displayed. You can review this information, but you don't need to do anything about it. Press a key to continue.
11. FIPS presents you with two numbers showing the size of the current Windows partition (on the left) and the size of the new Windows partition (on the right). (Figure 2-18 shows a sample screen.) Use the Up and Down arrow keys to adjust the space on each partition. Because you are changing the point at which the partition will be split in two, one number goes up as the other goes down. The number on the right (shown in MB) should be large enough for your Linux partition and swap partition.

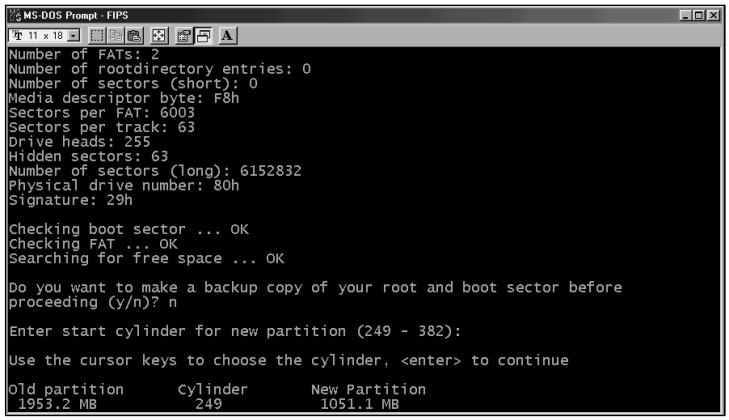

Figure 2-18 FIPS during the process of resizing the Windows partition

FIPS will not allow you to reduce the original Windows partition to a smaller size than is needed for the data residing on that Windows partition. If the number on the left for the size of the original Windows partition will not go low enough to allow space for Linux when the second Windows partition is deleted, you have not effectively moved all the data to the front of the Windows partition. Or the Windows partition may simply be too full to store Linux on the same computer.

12. Press **Enter** to accept your settings.
13. The partition table is displayed again as it will appear when updated by FIPS. If you are comfortable with the sizes shown, press **C** to confirm that you want to write this information to the hard disk. You can press **Ctrl+C** to exit FIPS without making any changes to your hard disk.
14. Reboot your computer immediately so that the updated hard disk information is reread into your system.

After using FIPS, you should have a smaller Windows partition containing the Windows operating system and your Windows data; you also have a second Windows partition that contains no data. During the Linux installation you will use a tool such as `fdisk` or the Disk Druid (discussed in Chapter 3) to delete this second, empty Windows partition and turn the resulting free space on the hard disk into the Linux partition and Linux swap partition.

In addition to FIPS, several commercial programs are available to help you prepare partitions for Linux. The most popular is called PartitionMagic, from PowerQuest (*www.powerquest.com*). PartitionMagic provides a graphical interface to create new partitions, and helps you back up data before altering the hard disk. However, PartitionMagic is commercial software that costs money. (The Caldera OpenLinux distribution includes a reduced version of PartitionMagic to help prepare a Windows system for a Linux installation.)

CHAPTER SUMMARY

- Many Linux installation programs are designed to detect a computer's hardware. However, an understanding of basic computer hardware terminology will help you install Linux more smoothly, because it also allows you to compile additional information that may be needed to complete the Linux installation.
- To use Linux networking, you must assign an IP address to the system. Additional IP addresses give Linux the ability to work with other computers on the network and to convert domain names into IP addresses.
- You can compile a system inventory by referring to the computer documentation, a vendor's Web site, or by checking the settings in a copy of Windows that is already installed on the computer. Networking information is assigned by a system administrator or ISP.
- Before installing Linux, you must prepare the hard disk by creating partition that will hold the Linux file system. You also need to leave room for a swap partition. You can use the FIPS utility to create an empty Windows partition on a hard disk dedicated to Windows. You can then delete the empty partition and use the space for Linux.

KEY TERMS

active partition — The partition that receives control from the BIOS when the system is turned on.

BIOS (Basic Input/Output System) — Information stored in ROM that provides instructions to the operating system for using the devices on a computer.

block — A unit of hard disk space that contains 1024 bytes, or two sectors.

boot manager — A program that lets you select an operating system to use each time you boot the computer.

boot record — A small area on each partition that contains a program to launch the operating system on that partition.

broadcast address — A special IP address that sends a packet of data to all computers on the local network.

byte — Space within a computer system sufficient to store one character.

cylinder — A set of tracks at the same location on all the platters of a hard disk.

defragmenting — A procedure that arranges each file in Windows so that the parts of the file are next to each other on the hard disk (as opposed to the parts of a file being fragmented, or spread across the hard disk).

direct memory access (DMA) channel — A communication method within a computer that allows a device to read and write directly to the computer's RAM, without going through the microprocessor first.

DNS server — A computer that uses the DNS protocol to convert from domain name and hostname to IP address.

domain name — A name assigned to a network.

Domain Name Service (DNS) — A protocol that maps human-readable domain names and hostnames to IP addresses that correspond to networks and individual computers.

dual-boot system — A computer that allows a person to choose which operating system to start each time the computer is booted (turned on).

Dynamic Host Configuration Protocol (DHCP) — A protocol that allows a computer to obtain an IP address dynamically from a network server at the time the computer is turned on.

ext2 — The default file system type for Linux.

FAT32 — The default file system type for Windows 98.

file system — An organized arrangement of information on a device such as a hard disk.

FIPS — A program that creates two separate partitions from an existing Windows partition.

fully qualified domain name (FQDN) — The hostname of a computer and the domain name of the network to which the computer is attached.

gateway address — The IP address of the computer on a local network that can send packets of data outside that network.

gigabyte (GB) — A measure of space on computers equal to 1024 megabytes, or roughly enough space to store 1 billion characters.

hard disk — Magnetic storage space for data such as the operating system and data files.

hexadecimal (hex) — A numbering system using base-16. Hex uses 0 to 9, plus the letters *A* through *F* (usually capitalized) to count the numbers 10 through 15.

host — A computer attached to a network.

hostname — A single word used to name a computer.

IDE — A low-cost, easy-to-manage interface used by most new computers to connect hard disks and CD-ROM drives to the CPU.

IDE controller — A computer hardware component used to communicate between an IDE-compatible hard disk or other IDE device and the microprocessor.

interrupt request (IRQ) — A numbered channel of communication allowing a device to inform the system that some action needs to be taken for the device.

I/O ports — Memory addresses used by a device for memory-mapped I/O.

IP — A networking protocol used to send packets of information across a network connection.

IP address — An identifying number assigned to a computer or device that uses IP to communicate across a network.

LILO (Linux Loader) — The boot manager included with Linux.

Master Boot Record (MBR) — A small area of the first hard disk that contains a program to decide how to start an operating system.

megabyte (MB) — A measure of space on computers equal to 1,048,576 bytes, or enough space to store roughly 1 million characters.

memory-mapped input/output (memory-mapped I/O) — A technique that assigns a range of memory addresses in a computer as a place for a device to send and receive data.

network mask — A set of numbers that tells the networking system in Linux how to identify IP addresses that are part of the local network.

NTFS — The default file system type for Windows NT.

partition — A distinct area of a hard disk that has been prepared to store a particular type of data.

protocol — An organized pattern of signals or words used to communicate efficiently.

random access memory (RAM) — Volatile electronic storage within a computer.

read-only memory (ROM) — Nonvolatile electronic storage within a computer. Used to store information about how the computer starts and how the devices in the computer are configured.

SCSI — A high-performance interface used by many types of devices to connect to a computer.

sector — A unit of data storage on a hard disk. Normally a sector contains 512 bytes.

swap partition (also called swap space) — A designated area on a hard disk used as virtual memory by the Linux kernel.

tracks — The concentric circles on each platter of a hard disk.

virtual memory — A feature of the Linux kernel that allows Linux to treat a special area of the hard disk as if it were RAM, storing programs there temporarily when they are not being used.

XFree86 Project — A free software project that creates software to provide X Window System functionality to Linux.

REVIEW QUESTIONS

1. You can purchase computer systems with Linux preinstalled. True or False?
2. Why is Linux sometimes considered more difficult to install than other operating systems?

3. A byte is enough space to store:
   a. One character
   b. One hexadecimal digit
   c. One sector
   d. One megabyte
4. When you see a value in gigabytes, it probably refers to:
   a. An IRQ number
   b. A monitor refresh rate
   c. Hard disk size
   d. RAM size
5. Which of the following cannot be changed by a user?
   a. Magnetic data on a hard disk
   b. Electronic storage in RAM
   c. The configuration settings in the BIOS
   d. The BIOS stored in ROM
6. The BIOS configuration settings can be accessed using a special menu system provided with most computers. True or False?
7. Which of the following is *not* part of the configuration for most devices on a computer?
   a. IRQ
   b. FIPS
   c. I/O ports
   d. DMA channel
8. The _____ numbering system is often used to refer to information about computer hardware.
   a. hexadecimal
   b. MBR
   c. binary
   d. SCSI
9. Which of the following is *not* a valid hexadecimal digit?
   a. E
   b. D
   c. A
   d. H

10. The _____ interface is a high-cost, high-performance method of connecting hard disks to a computer.
    a. LILO
    b. SCSI
    c. IDE
    d. MBR
11. An IP address consists of four numbers from 0 to 999, each separated by a period. True or False?
12. Name four items of information about the network that are required to complete a Linux installation, and explain the purpose of each one.
13. By using DHCP, a Linux system can:
    a. Convert a domain name to an IP address
    b. Pass packets outside the local network to the Internet
    c. Send broadcast messages to all computers on the network
    d. Obtain an IP address from a network server
14. An FQDN includes:
    a. A hostname and an IP address
    b. A hostname and a domain name
    c. A domain name and an IP address
    d. The IP address provided by DHCP
15. A DNS server provides the following service to computers on the network:
    a. It converts domain names to IP addresses.
    b. It provides IP addresses when a computer is turned on.
    c. It sends data packets to the Internet.
    d. It configures independent file systems on the hard disk.
16. The most important monitor specification to obtain before installing Linux is:
    a. Refresh rate (given in Hz)
    b. Scan rate (given in KHz)
    c. Cylinder count
    d. Frequency (given in MHz)
17. Name four methods of obtaining system hardware specifications prior to installing Linux.

18. The Windows _____ can provide many hardware details about a computer.
    a. Search dialog box
    b. FIPS program
    c. Start menu
    d. Control Panel
19. Networking information is not included in any Windows dialog boxes. True or False?
20. If you cannot find someone to assign an IP address to your computer, you can use any number temporarily. True or False?
21. A file system resides within a:
    a. Partition
    b. Track
    c. Block
    d. Sector
22. The BIOS has trouble booting Linux when:
    a. The boot record is on the active partition
    b. The Linux partition includes cylinder numbers beyond 1024
    c. The MBR refers to the active partition
    d. The IRQ of the BIOS is poorly chosen
23. Briefly define the purpose of a swap partition in Linux.
24. The FIPS program is used to split a Windows partition into two partitions. True or False?
25. When a file system is defragmented, the files are arranged so that:
    a. They are alphabetical in the directory structure.
    b. All parts of a file are next to each other on the hard disk.
    c. Compatible files are next to each other on the hard disk.
    d. Linux can easily use the files from Windows.

HANDS-ON PROJECTS

Project 2-1

In this activity you use the Windows Control Panel to find information about your hard disk. To complete this activity you need a computer with Windows installed. You can complete this activity as an exercise on a Windows computer even if you will not install Linux on the same computer later on.

1. Double-click the **My Computer** icon on the Windows Desktop. The My Computer window opens.
2. Double-click the **Control Panel** icon.
3. Double-click the **System** icon in the Control Panel window. The System Properties dialog box opens.
4. Click the **Device Manager** tab.
5. Locate the item in the Device Manager window list labeled Disk drives. Click the **+** to the left of the Disk drives label. Can you identify the interface used by the disk drives on your system based on what you see here? Does this dialog box appear to provide any information about how the partitions on the hard disk are set up?
6. Double-click the icon in the list representing a hard disk. The properties dialog box for the hard disk opens. Does the information in this dialog box help you learn about the partitions on the disk?
7. Click **Cancel** to close the hard disk properties dialog box.
8. Click **Cancel** to close the System Properties dialog box.

Project 2-2

In this activity you review the information provided on your Linux installation CD about the FIPS program. The documentation for FIPS includes information about how hard disks are organized and explains how to use FIPS. To complete this activity you need a CD containing a Linux distribution and a computer with a CD-ROM drive. This exercise assumes that you have a Red Hat Linux CD and a computer with Windows installed and a CD-ROM drive. Other types of Linux CDs can be used, but the location of the FIPS program on the CD will differ.

1. Insert your Linux CD into the CD-ROM drive.
2. Double-click the **My Computer** icon on the Windows Desktop.
3. Double-click the icon for the drive in which the Linux CD is located.
4. Double-click the **dosutils** icon.
5. Double-click the **fipsdocs** icon.
6. Double-click the **fips** document icon. (This icon may be labeled fips.doc, depending on your Windows configuration). The FIPS documentation appears in a word processing program.
7. Press **Alt+Tab** to switch back to the window showing the files in the `fipsdoc` directory.

8. Double-click on another document icon such as the techinfo (or techinfo.txt) icon.
9. Review the document files that you have opened to learn more about how FIPS operates. Open other files in the `fipsdoc` directory to learn additional information.

Project 2-3

In this activity you use the Windows Control Panel to gather information about how networking is set up on your computer. To complete this activity your computer must have Windows installed, and it must be connected to a network, either via a network adapter such as an Ethernet card, or via a modem. You can complete this activity as an exercise on a Windows computer even if you will not install Linux on the same computer later on.

1. Double-click the **My Computer** icon on the Windows Desktop. The My Computer window opens.
2. Double-click the **Control Panel** icon. The Control Panel window opens.
3. Double-click the **Network** icon.
4. Click the **Configuration** tab in the Network dialog box.
5. In the list of networking components in the Network dialog box, select the **TCP/IP** item.
6. Click **Properties** below the list of components. A TCP/IP Properties dialog box appears.
7. Click the **DNS Configuration** tab.
8. If Enable DNS is selected, write down the information in the Host, Domain, and DNS Server Search Order fields. Save this information to use when you install Linux.
9. Close all open dialog boxes without making any changes.

CASE PROJECTS

1. You have been asked to install Linux in a lab containing about 40 computers. Some of the computers are old; others are quite new. You have no experience with these particular machines, so you don't yet know anything about what they contain. Describe the steps you will take, in order, to quickly determine the different devices on which you must install Linux.
2. Assuming that some of the computers in the lab do not have any operating system installed, how does the process in Step 1 change? What resources outside of the lab might you draw on to help you identify the devices used in each computer? How might the vendors of the computers help you?
3. Now suppose you have finished studying the hardware in all of the lab's computers. Before you begin installing Linux, you have to decide which distribution of Linux to use. What would be the advantage of using the same Linux distribution on all the lab's computers? What factors might cause you to choose a different distribution, either before installation began or after you had begun to install Linux on multiple systems?